

Latent Sentiment Detection in Online Social Networks: A Communications-oriented View

Vinay Uday Prabhu^{1,2}, Rohit Negi¹ and Miguel Rodrigues³

¹Electrical and Computer Engineering Department, Carnegie Mellon University

²Departamento de Ciencia de Computadores, University of Porto

³Electronic and Electrical Engineering Department, University College London

Email:vinayprabhu@cmu.edu, negi@ece.cmu.edu, m.rodrigues@ucl.ac.uk

Abstract

In this paper, we consider the problem of latent sentiment detection in Online Social Networks such as Twitter. We demonstrate the benefits of using the underlying social network as an Ising prior to perform network-aided sentiment detection. We show that the use of the underlying network results in substantially lower detection error rates compared to strictly features-based detection. In doing so, we introduce a novel communications-oriented framework for characterizing the probability of error, based on information-theoretic analysis. We study the variation of the calculated error exponent for several stylized network topologies such as the complete network, the star network and the closed-chain network, and show the importance of the network structure in determining detection performance.

Index Terms

Sentiment detection, Online Social Networks, Error analysis.

I. INTRODUCTION

Online Social Networks (OSN), such as Twitter¹ have come to heavily influence the way people socially galvanize. Recent world events such as the Arab Spring, witnessed cascading democratic revolutions characterized by a strong reliance on online social media such as Twitter and Facebook [1]. Today there are about 554,750,000 registered active Twitter users with about

¹ <https://twitter.com/>

135,000 new Twitter users signing up everyday. Around 58 million tweets are *tweeted* per day and the website attracts over 190 million visitors every month [2]. Such staggering numbers have turned such OSN into a veritable data gold mine for organizations and individuals who have a strong social, political or economic interest in maintaining and enhancing their clout and reputation. Therefore, extracting and analyzing the embedded sentiment in the microblogs (or *Tweets*) posted by the tweeters about these organizations or individuals, or specific issues, products and events related to them or their competitors, is of great interest to them. Of particular interest is the *latent sentiment* (as opposed to individual expressed sentiments), which can be either positive or negative with respect to a particular position. We explain this latent sentiment in detail in Section II. Strict length restrictions (such as the 140 character-limit per tweet), irregular structure of the microblog content and the usage of sarcasm renders the problem of automatic latent sentiment detection (classifying latent sentiment as positive or negative) from the microblog contents error-prone. As evidenced in literature [3], [4], sentiment detection has been approached from an engineering perspective with the main focus being on sentiment detection algorithms, followed by empirical performance comparisons using standard datasets such as Stanford Twitter Sentiment (STS) dataset and the Obama-McCain Debate dataset [5]. We had considered in [6], network aided detection of votes in the senate harnessing the joint press release network. In this paper, we approach such problems from a relatively scientific perspective. That is, we attempt to answer the following two important questions regarding social-network aided latent sentiment detection.

- 1) How much benefit can be expected by using the social network for latent sentiment detection?
- 2) How does the social network structure affect the performance of network-aided sentiment detection?

Towards this end, we use a stylized model of latent sentiment detection, based on a Markov Random Field model. We then analyze the performance of the optimal Maximum A Posteriori Probability (MAP) sentiment detector, keeping in mind the underlying social network structure. For this, we are inspired by a communications-oriented viewpoint, where we view the underlying network as providing a *weak channel code*, that transmits one bit of information. Accordingly, we are able to analyze the performance of the sentiment detector, by borrowing tools from

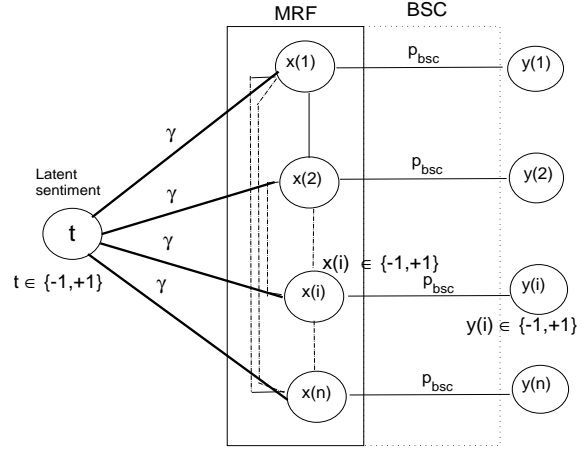


Fig. 1. Model for latent sentiment detection.

information theory. We can then compute and contrast the performance under various stylized social network topologies, thus providing answers to both the questions posed above. Thus, we show that communication theorists can contribute to the growing field of social network analysis.

The rest of the paper is organized as follows. In Section II, we describe the latent sentiment detection problem, introduce the formal model and motivate its relevance through real-world scenarios based on Twitter. We also specify the optimal MAP sentiment detector. In Section III, we perform a communications-oriented analysis of the MAP sentiment detector to derive an upper bound on the detection error probability, in terms of an error exponent. We also show how the exponent can be evaluated numerically for various stylized topologies such as the complete network, the star network and the chain network. In Section IV, we present numerical results that show how the error exponent depends on the network topology and other social network parameters. We conclude the paper in Section V.

II. LATENT SENTIMENT DETECTION PROBLEM

A. Model

The model considered in this paper is shown in Fig. 1. Let $\mathbf{x} \in \{-1, +1\}^n$ be the vector of expressed sentiments of the n members of a social network, with $x(i) \in \{-1, +1\}$ being the expressed sentiment of the i^{th} member/node. The social network structure is modeled as an undirected graph $G(V, E)$ characterized by its symmetric adjacency matrix A . It may be obtained using the follower/followee relationships, or in some cases, using the @-mentions in the tweets

[7]. The graph is undirected since we will use it to model correlation, rather than influence flows. The sentiments are assumed to be sampled from an underlying homogeneous Markov Random Field (MRF) [8] with unit edge potential and inverse temperature parameter, θ . In this paper, we assume $\theta \geq 0$, so that we are restricting ourselves to attractive/ferromagnetic models, which correspond to homophilic networks. In such a ferromagnetic model, the neighboring nodes positively correlate with each other, so that the distribution has more probability in configurations with similar values on the nodes of the graph. Let $t \in \{-1, +1\}$ indicate the latent sentiment variable which homogeneously influences every node of the network as a local field of strength γt . In the absence of any sentiment bias, we assume t to be equi-probably equal to $+1$ or -1 . Thus, the conditional distribution of \mathbf{x} given t , can be written as,

$$p(\mathbf{x}|t) = \frac{\exp \{ \theta \mathbf{x}^T A \mathbf{x} + \gamma t \mathbf{e}^T \mathbf{x} \}}{\sum_{\mathbf{x}} \exp \{ \theta \mathbf{x}^T A \mathbf{x} + \gamma t \mathbf{e}^T \mathbf{x} \}}. \quad (1)$$

Notice that from the communications perspective, \mathbf{x} is a codeword randomly chosen in response to bit t . Let \mathbf{y} be a noisy estimate of \mathbf{x} . It may be estimated from the features extracted from the user profiles or could even be the sentiment vector estimated by a given classifier algorithm, such as the ones in [3] and [4]. While the alphabet of each $y(i)$ can be arbitrary, in this paper, for simplicity, we assume that it is binary $\{-1, +1\}$. We model \mathbf{y} to be the output of n -identical and independent Binary Symmetric Channels (BSCs) characterized by the equal cross-over probability p_{bsc} , with the elements of the true sentiment vector \mathbf{x} as the input. Therefore, the conditional distribution, $p(\mathbf{y}|\mathbf{x})$ may be written as,

$$p(\mathbf{y}|\mathbf{x}) = \frac{1}{c^n} \exp \{ \varepsilon \mathbf{y}^T \mathbf{x} \}, \quad (2)$$

where $\varepsilon = \frac{1}{2} \log \left(\frac{p_{bsc}}{1-p_{bsc}} \right)$ and $c = 2 \cosh(\varepsilon)$. The joint distribution of all variables may now be written as,

$$p(t, \mathbf{x}, \mathbf{y}) = \frac{1}{2 Z(t)} \exp \{ \theta \mathbf{x}^T A \mathbf{x} + \varepsilon \mathbf{y}^T \mathbf{x} + \gamma t \mathbf{e}^T \mathbf{x} \}, \quad (3)$$

$$Z(t) = c^n \sum_{\mathbf{x}} \exp \{ \theta \mathbf{x}^T A \mathbf{x} + \gamma t \mathbf{e}^T \mathbf{x} \}. \quad (4)$$

B. Real-world examples from Twitter

The model (3) is applicable to several real-world latent sentiment detection scenarios. We begin by assuming that there exists a latent sentiment ($t \in \{-1, +1\}$), which will cause a

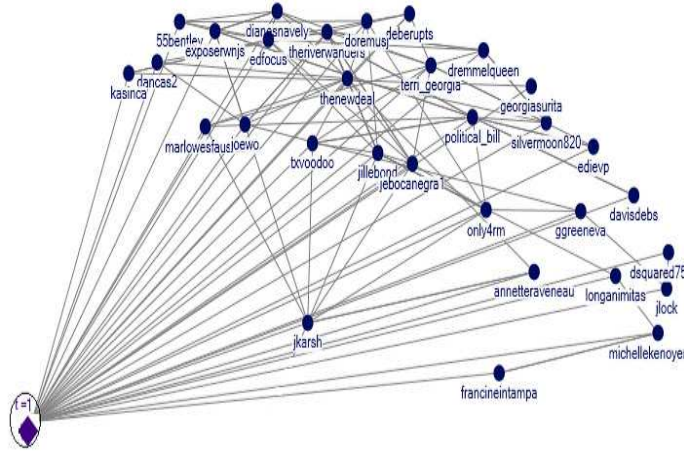


Fig. 2. The #iloveobamacare network of sentiments.

certain concrete event in the future. This event may be the passage (or defeat) of a bill in the senate, or an up (or down) movement of the stock market, when $t = +1$ (or -1 , respectively). The intention is to predict this real-world event in the present using the expressed sentiments gathered from the *twitterverse* (y in our model). Thus, it is the same as detecting the value of t (hence the term *latent sentiment detection*).

Fig. 2 and Fig. 3 jointly refer to one such scenario where the underlying sentiment t was support (+1) or opposition (-1) to the Patient Protection and Affordable Care Act (PPACA) [9], nicknamed as *Obamacare*. Fig. 2 represents the social network of liberal-minded follower/followee networks of twitter users, who tweeted in support of the #iloveobamacare hashtag, which was promulgated on twitter to galvanize more support. Fig. 3 represents the follower/followee network of conservative-minded twitter users who attacked the #iloveobamacare hashtag with a series of sharp and sarcastic tweets resulting in what is called **Hashtag-Hijacking** [10]. The national survey conducted by the Pew Research Center and USA TODAY [11], later confirmed the underlying sentiment of support ($t = +1$) to the act amongst liberals and opposition ($t = -1$) to the act amongst the conservatives.

The goal would be for an automatic sentiment detector to predict t for each network. Of course, in this example, one can do this knowing the political stance of the networks, which is side information. However, automatic detection aims to apply a general method based detector on y without requiring human intervention through specialized side information.

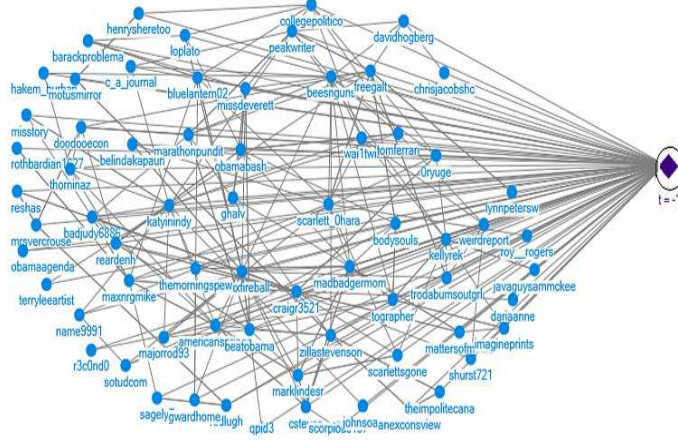


Fig. 3. The 'anti-supportive' #iloveobamacare network $-t = -1$ scenario

C. MAP sentiment detector

We assume that the network adjacency matrix A and the other system parameters, θ, ϵ, γ are known. Then, the optimal MAP latent sentiment detector \hat{t} , which is equal to the Maximum Likelihood (ML) detector in this case, is

$$\hat{t} = \begin{cases} +1, & l(\mathbf{y}) \geq 1 \\ -1, & l(\mathbf{y}) < 1 \end{cases} \quad (5)$$

where, the likelihood $l(\mathbf{y})$ is,

$$l(\mathbf{y}) = \frac{\sum_{\mathbf{x}} p(t = 1, \mathbf{x}, \mathbf{y})}{\sum_{\mathbf{x}} p(t = -1, \mathbf{x}, \mathbf{y})}. \quad (6)$$

In the next section, we analyze the performance of the optimal MAP sentiment detector (5). We perform a communications-inspired analysis of the probability of error of the detector, using which, we seek to understand the contribution of the social network, as well as the role played by the underlying network topology in the performance of the detector.

III. COMMUNICATIONS-INSPIRED ANALYSIS OF SENTIMENT DETECTION

In this section, we perform an analysis of the error probability of the latent sentiment detector (5). By symmetry of the model, the error probability is,

$$\begin{aligned} P_e &= P_{e|t=-1} \\ &= P(l(\mathbf{y}) \geq 1 | t = -1) \end{aligned}$$

$$= \sum_{\mathbf{y}} p(\mathbf{y}|t = -1) \mathbf{1}(l(\mathbf{y}) > 1). \quad (7)$$

Here, $\mathbf{1}(A) = 1$ for the event A , otherwise it is 0. P_e is infeasible to calculate for large social network. So, in the next subsection, we present an upper bound for P_e .

A. P_e upper bound

The main result of this paper is the following theorem. *Theorem 1:* For the optimal MAP detector (5), an upper bound on the error probability P_e is,

$$\begin{aligned} P_{e,UB} &= \frac{1}{Z(\theta, \gamma)(\cosh(\varepsilon))^n} \min_b A(b), \quad \text{where,} \\ A(b) &= \left(\frac{\cosh(2b) + \cosh(2\varepsilon)}{2} \right)^{n/2} Z(\theta, \beta), \\ Z(\theta, \beta) &= \sum_{\mathbf{x}} \exp \{ \theta \mathbf{x}^T \mathbf{A} \mathbf{x} - \beta \mathbf{e}^T \mathbf{x} \}, \\ Z(\theta, \gamma) &= \sum_{\mathbf{x}} \exp \{ \theta \mathbf{x}^T \mathbf{A} \mathbf{x} - \gamma \mathbf{e}^T \mathbf{x} \}. \end{aligned} \quad (8)$$

Proof: The proof relies on Information-theoretic analysis. (See Appendix).

B. Computation of the upper bound

From (8), we see that computation of the upper bound requires computing the partition functions related to the underlying MRF. This is an NP-complete problem in general [12]. However, for certain stylized topologies, such as the complete network, the star network and the (closed) chain network and others, there exist closed form expressions for the partition function. We compute the partition functions in closed form for these topologies and list them below.

$$\begin{aligned} Z_{\text{complete}}(\theta, \gamma) &= \sum_{m=0}^n \binom{n}{m} u_c^{(n-2m)^2 - n} v_c^{(n-2m)} \\ u_c &= \exp\left(\frac{\theta}{2}\right), \quad v_c = \exp(\gamma). \\ Z_{\text{chain}}(\theta, \gamma) &= \lambda_+^n + \lambda_-^n \\ \lambda_{\pm}^n &= \exp(\theta) \left\{ \cosh(\gamma) \pm \sqrt{\sinh^2(\gamma) + \exp(-4\theta)} \right\}. \end{aligned} \quad (9)$$

$$\begin{aligned}
Z_{star}(\theta, \gamma) &= u^{n-1} \left[v^n \sum_{m=0}^{n-1} \binom{n-1}{m} (uv)^{-2m} \right. \\
&\quad \left. + v^{-n} \sum_{k=0}^{n-1} \binom{n-1}{k} \left(\frac{u}{v}\right)^{-2k} \right] \\
u &= \exp(\theta), v = \exp(\gamma).
\end{aligned} \tag{10}$$

Thus the error probability upper bound (8) can be evaluated for the above network topologies, to provide insight into the impact of social network structure on the sentiment detector performance. Note that, via (8), we have reduced the complicated problem of computing an error probability to a problem of calculating an MRF partition function. The partition function calculation is a well researched problem in MRF theory [13], [14] [15], and significant effort has been expended in statistical physics and machine learning to compute it for a variety of graphs. Thus, our theorem facilitates importing ideas from that literature to obtain the error probability bound for a variety of graphs.

C. Error exponent

For large networks, an error exponent can be defined as,

$$\alpha = \lim_{n \rightarrow \infty} \inf \left\{ \frac{-\log P_e}{n} \right\}. \tag{11}$$

Using the bound (8), we can show that,

$$\begin{aligned}
\alpha &\geq \log(\cosh(\varepsilon)) + \frac{\log(Z(\theta, \gamma))}{n} - \\
&\min_b \left[\frac{1}{2} \log \left(\frac{\cosh(2b) + \cosh(2\varepsilon)}{2} \right) + \frac{\log(Z(\theta, \beta))}{n} \right].
\end{aligned} \tag{12}$$

For the simple case where the network is absent, α can be shown to be exactly,

$$\begin{aligned}
\alpha_{iid} &= \log(\cosh(\varepsilon)) + \log(\cosh(\gamma)) \\
&\quad - \frac{1}{2} \log(\cosh^2(\varepsilon) + \cosh^2(\gamma)).
\end{aligned} \tag{13}$$

Thus, we can use (13) as a benchmark to compare error exponents obtained for various network topologies via evaluation of (12). In the following section, we perform these comparisons by plotting the variation of the error exponent α derived in (12) and (13) with respect to the model parameters, θ , γ and ε .

IV. NUMERICAL RESULTS: ERROR EXPONENT OF DIFFERENT NETWORKS

The aim of this section is to answer the two questions raised in the Introduction. First, to show how much improvement can be expected when the network is used for detection and secondly, to understand the effect of network topology on MAP detector performance. We do this by plotting (12).

The inverse temperature parameter, θ , captures how strongly opinions are correlated in the network. Increasing θ increases the homophily in the OSN, which makes the underlying MRF more strongly ferromagnetic. For example, nodes of a particular political party will be strongly homophilic. This is exploited by the network aided sentiment detector in (5), resulting in faster increase of α for network aided detectors with increasing θ , as shown in Fig. 4, with $\varepsilon = 0.5$ and $\gamma = 0.5$. As seen, α for the no-network (i.i.d.) scenario remains invariant to the change in θ while for the network-aided cases, we see $\alpha \rightarrow \log(\cosh(\varepsilon)) = 0.1201$ as θ increases. The rate of increase of α is clearly the highest for the complete network case when compared to the star and chain networks. Notice also that the complete network reaches the maximum α beyond a certain threshold θ since it is strongly homophilic.

The parameter γ captures how strongly the latent sentiment influences the opinions of the users/nodes. Practically, large γ pertains to a highly emotive situation where the users are strongly influenced to take a particular position even in the absence of strong homophily (θ). Since γ is independent of the network effect, both the No-net and the network-aided detectors can harness the effect of increase in γ resulting in higher α in both scenarios. However, the combined use of the network as well as γ results in higher α for the network aided cases. This behavior is seen in Fig. 5. The other system parameters were held constant ($\varepsilon = 0.5$ and $\theta = 0.5$). Again, $\alpha \rightarrow \log(\cosh(\varepsilon)) = 0.1201$ as γ increases similar to Fig. 4.

Finally, we turn our attention towards Fig. 6 which shows the variation of α with ε for $\theta = 0.5$ and $\gamma = 0.5$. $-\varepsilon$ captures the amount of noise in the BSC in dB (with $\varepsilon \rightarrow \infty$ being the zero-noise case). Alternately, ε is the accuracy of the detector used to obtain estimated individual sentiments, \mathbf{y} . As expected, increase in ε results in higher α for both network-aided and the no-network cases. In all cases, the error exponent α of the complete network is significantly larger than the star and the chain networks.




Fig. 4. Variation of error error exponent (α) with θ for different network topologies




Fig. 5. Variation of the error exponent (α) with γ for different network topologies

V. CONCLUSION

In this paper, we have introduced a novel communications-inspired framework for analyzing probability of error of network-aided detection of latent sentiment in Online Social Networks. Through this, we have attempted to provide insight into the role played by the network, specifically the topology, in lowering the probability of error of detection, thereby rigorously characterizing the worth of the network as an additional information source for sentiment detection. Firstly, we motivate the practical scenarios where the model is applicable and then provide an analysis of the upper bound on the probability of error, or equivalently, the error exponent




Fig. 6. Variation of the error exponent (α) with ϵ for different network topologies

for large networks. Finally, we plot the variation of this upper bound with respect to model parameters for the complete network, star network and chain network topologies, and show the improvement in performance relative to non-network-aided detection.

APPENDIX

Following [16], we can upper bound the probability of error using an arbitrary function $0 \leq s(\mathbf{y}) \leq 1$ as,

$$\begin{aligned}
 P_e &\leq \sum_{\mathbf{y}} p(\mathbf{y}|t = -1) (l(\mathbf{y}))^{s(\mathbf{y})} \\
 &= \sum_{\mathbf{y}} p(\mathbf{y}|t = -1)^{(1-s(\mathbf{y}))} p(\mathbf{y}|t = +1)^{s(\mathbf{y})} \\
 &\leq \sum_{\mathbf{y}} 2 \times \left(\sum_{\mathbf{x}} p(t = -1, \mathbf{x}, \mathbf{y}) \right)^{(1-s(\mathbf{y}))} \\
 &\quad \times \left(\sum_{\mathbf{x}} p(t = +1, \mathbf{x}, \mathbf{y}) \right)^{s(\mathbf{y})} \\
 &\leq 2 \sum_{\mathbf{y}} \left(\sum_{\mathbf{x}} \frac{\exp \{ \theta \mathbf{x}^T A \mathbf{x} + \epsilon \mathbf{y}^T \mathbf{x} - \gamma \mathbf{e}^T \mathbf{x} \}}{Z(-1)} \right)^{1-s(\mathbf{y})} \\
 &\quad \times \left(\sum_{\mathbf{x}} \frac{\exp \{ \theta \mathbf{x}^T A \mathbf{x} + \epsilon \mathbf{y}^T \mathbf{x} + \gamma \mathbf{e}^T \mathbf{x} \}}{Z(+1)} \right)^{s(\mathbf{y})}.
 \end{aligned}$$

By optimizing over the function $s(\mathbf{y})$, it is possible to write a tighter upper bound,

$$P_e \leq \frac{1}{c^n} \sum_{\mathbf{y}} \min_{s(\mathbf{y})} \left(E \left(\exp \{ +\varepsilon \mathbf{y}^T \mathbf{x} \} \right) \right)^{(1-s(\mathbf{y}))} \times \\ \left(E \left(\exp \{ -\varepsilon \mathbf{y}^T \mathbf{x} \} \right) \right)^{s(\mathbf{y})}$$

by using (4). Here, $E(\cdot)$ denotes expectation with respect to \mathbf{x} , conditioned on $t = -1$ (see (1)).

Using

$$\min_{0 \leq s \leq 1} \left(a_1^s a_2^{(1-s)} \right) \leq \frac{a_1 e^s + a_2 e^{-s}}{e^s + e^{-s}} \leq \frac{a_1 e^s + a_2 e^{-s}}{2},$$

we continue the upper bound as,

$$P_e \leq \frac{1}{2c^n} \sum_{\mathbf{y}} \min_{s(\mathbf{y})} \left\{ \left(E \left(\exp \{ +\varepsilon \mathbf{y}^T \mathbf{x} \} \right) \right)^{-s(\mathbf{y})} + \right. \\ \left. \left(E \left(\exp \{ -\varepsilon \mathbf{y}^T \mathbf{x} \} \right) \right)^{s(\mathbf{y})} \right\}$$

and by writing $s(\mathbf{y}) = -\mathbf{b}^T \mathbf{y}$, we obtain

$$P_e \leq \frac{1}{2c^n} \sum_{\mathbf{y}} \min_{\mathbf{b}} \left\{ E \left(\exp \{ -\mathbf{y}^T (\mathbf{b} + \varepsilon \mathbf{x}) \} \right) + \right. \\ \left. E \left(\exp \{ \mathbf{y}^T (\mathbf{b} + \varepsilon \mathbf{x}) \} \right) \right\}$$

By observing that

$$\sum_{\mathbf{y}} E \left(\exp \{ -\mathbf{y}^T (\mathbf{b} + \varepsilon \mathbf{x}) \} \right) + E \left(\exp \{ \mathbf{y}^T (\mathbf{b} + \varepsilon \mathbf{x}) \} \right) = \\ 2 \left(\prod_{i=1}^n \exp \{ b_i + \varepsilon x_i \} + \exp \{ -b_i + \varepsilon x_i \} \right)$$

and interchanging the order of summation and minimization, we simplify the upper bound,

$$P_e \leq \frac{1}{\left(\frac{c}{2}\right)^n} \min_{\mathbf{b}} E \left(\prod_{i=1}^n \cosh(b_i + \varepsilon x_i) \right). \quad (14)$$

Now, setting $\mathbf{b} = \mathbf{b}\mathbf{e}$ (which would loosen the bound) and using the identity $\prod_{i=1}^n \cosh(b + \varepsilon x_i) =$

$\left(\frac{\cosh(2b) + \cosh(2\varepsilon)}{2} \right)^{n/2} \left(\sqrt{\frac{\cosh(b+\varepsilon)}{\cosh(b-\varepsilon)}} \right)^{\mathbf{e}^T \mathbf{x}}$, the upper bound simplifies further to

$$P_e \leq \frac{\min_b E \left(\left(\frac{\cosh(2b) + \cosh(2\varepsilon)}{2} \right)^{n/2} \left(\sqrt{\frac{\cosh(b+\varepsilon)}{\cosh(b-\varepsilon)}} \right)^{\mathbf{e}^T \mathbf{x}} \right)}{(\cosh(\varepsilon))^n}.$$

We can finally expand the expectation within the minimization to produce the upper bound:

$$P_{e,UB} = \frac{1}{(\cosh(\varepsilon))^n} \min_b \left(\left(\frac{\cosh(2b) + \cosh(2\varepsilon)}{2} \right)^{n/2} \times \frac{\sum_{\mathbf{x}} \exp\{\theta \mathbf{x}^T A \mathbf{x} - \beta \mathbf{e}^T \mathbf{x}\}}{\sum_{\mathbf{x}} \exp\{\theta \mathbf{x}^T A \mathbf{x} - \gamma \mathbf{e}^T \mathbf{x}\}} \right)$$

or, by recognizing that the summations (over \mathbf{x}) have the form of the standard MRF partition function, we finally obtain the upper bound,

$$P_{e,UB} = \frac{1}{Z(\theta, \gamma)(\cosh(\varepsilon))^n} \min_b \left(\left(\frac{\cosh(2b) + \cosh(2\varepsilon)}{2} \right)^{n/2} \times Z(\theta, \beta) \right),$$

where $\beta = \gamma + \frac{1}{2} \log \left(\frac{\cosh(b-\varepsilon)}{\cosh(b+\varepsilon)} \right)$, $Z(\theta, \beta) = \sum_{\mathbf{x}} \exp\{\theta \mathbf{x}^T A \mathbf{x} - \beta \mathbf{e}^T \mathbf{x}\}$ and $Z(\theta, \gamma) = \sum_{\mathbf{x}} \exp\{\theta \mathbf{x}^T A \mathbf{x} - \gamma \mathbf{e}^T \mathbf{x}\}$.

REFERENCES

- [1] P. N. Howard and M. M. Hussain, *Democracy's Fourth Wave?: Digital Media and the Arab Spring*. Oxford University Press, 2013.
- [2] [Online]. Available: <http://www.statisticbrain.com/twitter-statistics/>
- [3] L. Barbosa and J. Feng, "Robust sentiment detection on twitter from biased and noisy data," in *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*. Association for Computational Linguistics, 2010, pp. 36–44.
- [4] E. Kouloumpis, T. Wilson, and J. Moore, "Twitter sentiment analysis: The good the bad and the omg!" in *ICWSM*, 2011.
- [5] H. Saif, Y. He, and H. Alani, "Semantic sentiment analysis of twitter," in *The Semantic Web—ISWC 2012*. Springer, 2012, pp. 508–524.
- [6] V. Prabhu, R. Negi, and M. Rodrigues, "Bipartisan cloture roll call vote prediction using the joint press release network in us senate," in *ICML workshop on Structured Learning: (SLG 2013)*, 2013.
- [7] C. Tan, L. Lee, J. Tang, L. Jiang, M. Zhou, and P. Li, "User-level sentiment analysis incorporating social networks," in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2011, pp. 1397–1405.
- [8] K. Binder, "Ising model," in *Encyclopedia of Mathematics*. Springer Publishers, 2001, iISBN 9781-55608.
- [9] [Online]. Available: <http://www.gpo.gov/fdsys/pkg/PLAW-111publ148/html/PLAW-111publ148.htm>
- [10] [Online]. Available: <http://dailycaller.com/2012/03/23/conservatives-hijack-ilikeobamacare-hashtag-on-twitter/>
- [11] [Online]. Available: <http://www.people-press.org/2013/09/16/as-health-care-law-proceeds-opposition-and-uncertainty-persist/>
- [12] B. A. Cipra, "The ising model is np-complete," *SIAM News*, vol. 33, no. 6, pp. 07–00, 2000.
- [13] P. W. Kasteleyn, "Dimer statistics and phase transitions," *Journal of Mathematical Physics*, vol. 4, p. 287, 1963.
- [14] N. N. Schraudolph and D. Kamenetsky, "Efficient exact inference in planar ising models," *arXiv:0810.4401*, 2008.
- [15] B. Flach, "A class of random fields on complete graphs with tractable partition function," *arXiv:1212.2136*, 2012.
- [16] R. Gallager, "A simple derivation of the coding theorem and some applications," *Information Theory, IEEE Transactions on*, vol. 11, no. 1, pp. 3–18, 1965.